

100% Pass 2026 Efficient DP-600: Implementing Analytics Solutions Using Microsoft Fabric Reliable Test Testking



2026 Latest Actual4Exams DP-600 PDF Dumps and DP-600 Exam Engine Free Share: https://drive.google.com/open?id=1kAgKVATMZ07TySwH2gT5t4UbIMxv_5-6

Perhaps the path to successful pass the DP-600 is filled variables, but now there is only one possibility to successfully obtain a DP-600 certification. That is to download and use our DP-600 study materials. Trying to become a DP-600 certified professional. Then join our preparation kit. DP-600 is an excellent platform that provides an DP-600 study materials that are officially equipped by an expert. Our DP-600 Exam Material can be studied and passed quickly within one week of the exam. Our DP-600 exam materials will give you the best knowledge of the contents of the DP-600 exam certification course outline. Our DP-600 materials provide you with the best learning prospects and give you more than you expect by adopting minimal effort.

Microsoft DP-600 Exam Syllabus Topics:

Topic	Details
Topic 1	<ul style="list-style-type: none"> Prepare data: This section of the exam measures the skills of engineers and covers essential data preparation tasks. It includes establishing data connections and discovering sources through tools like the OneLake data hub and the real-time hub. Candidates must demonstrate knowledge of selecting the appropriate storage type—lakehouse, warehouse, or eventhouse—depending on the use case. It also includes implementing OneLake integrations with Eventhouse and semantic models. The transformation part involves creating views, stored procedures, and functions, as well as enriching, merging, denormalizing, and aggregating data. Engineers are also expected to handle data quality issues like duplicates, missing values, and nulls, along with converting data types and filtering. Furthermore, querying and analyzing data using tools like SQL, KQL, and the Visual Query Editor is tested in this domain.
Topic 2	<ul style="list-style-type: none"> Maintain a data analytics solution: This section of the exam measures the skills of administrators and covers tasks related to enforcing security and managing the Power BI environment. It involves setting up access controls at both workspace and item levels, ensuring appropriate permissions for users and groups. Row-level, column-level, object-level, and file-level access controls are also included, alongside the application of sensitivity labels to classify data securely. This section also tests the ability to endorse Power BI items for organizational use and oversee the complete development lifecycle of analytics assets by configuring version control, managing Power BI Desktop projects, setting up deployment pipelines, assessing downstream impacts from various data assets, and handling semantic model deployments using XMLA endpoint. Reusable asset management is also a part of this domain.

Topic 3	<ul style="list-style-type: none"> • Implement and manage semantic models: This section of the exam measures the skills of architects and focuses on designing and optimizing semantic models to support enterprise-scale analytics. It evaluates understanding of storage modes and implementing star schemas and complex relationships, such as bridge tables and many-to-many joins. Architects must write DAX-based calculations using variables, iterators, and filtering techniques. The use of calculation groups, dynamic format strings, and field parameters is included. The section also includes configuring large semantic models and designing composite models. For optimization, candidates are expected to improve report visual and DAX performance, configure Direct Lake behaviors, and implement incremental refresh strategies effectively.
---------	--

>> DP-600 Reliable Test Testking <<

DP-600 Latest Study Guide, New DP-600 Practice Materials

Our DP-600 exam questions are famous for the good performance and stable operation. Customers usually attach great importance on the function of a product. So after a long period of research and development, our DP-600 learning prep has been optimized greatly. We can promise that all of your operation is totally flexible. Even if we come across much technology problems, we have never given up. Also, we take our customers' suggestions of the DP-600 Actual Test guide seriously. Sometimes, we will receive some good suggestions from our users. Once our researchers regard it possible to realize, we will try our best to perfect the details of the DP-600 learning prep. We are keeping advancing with you. You will regret if you do not choose our study materials.

Microsoft Implementing Analytics Solutions Using Microsoft Fabric Sample Questions (Q65-Q70):

NEW QUESTION # 65

You are analyzing customer purchases in a Fabric notebook by using PySpark. You have the following DataFrames:

□ You need to join the DataFrames on the customer_id column. The solution must minimize data shuffling. You write the following code.

□ Which code should you run to populate the results DataFrame?

- A. □
- **B.** □
- C. □
- D. □

Answer: B

Explanation:

The correct code to populate the results DataFrame with minimal data shuffling is Option A.

Using the broadcast function in PySpark is a way to minimize data movement by broadcasting the smaller DataFrame (customers) to each node in the cluster.

This is ideal when one DataFrame is much smaller than the other, as in this case with customers.

References = You can refer to the official Apache Spark documentation for more details on joins and the broadcast hint.

NEW QUESTION # 66

You need to create a DAX measure to calculate the average overall satisfaction score.

How should you complete the DAX code? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

□

Answer:

Explanation:

□ Explanation:

* The measure should use the AVERAGE function to calculate the average value.

* It should reference the Response Value column from the 'Survey' table.

* The 'Number of months' should be used to define the period for the average calculation.

To calculate the average overall satisfaction score using DAX, you would need to use the AVERAGE function on the response values related to satisfaction questions. The DATESINPERIOD function will help in calculating the rolling average over the last 12 months.

NEW QUESTION # 67

Case Study 2 - Litware, Inc

Overview

Litware, Inc. is a manufacturing company that has offices throughout North America. The analytics team at Litware contains data engineers, analytics engineers, data analysts, and data scientists.

Existing Environment

Fabric Environment

Litware has been using a Microsoft Power BI tenant for three years. Litware has NOT enabled any Fabric capacities and features.

Available Data

Litware has data that must be analyzed as shown in the following table.

The Product data contains a single table and the following columns.

The customer satisfaction data contains the following tables:

- Survey
- Question
- Response

For each survey submitted, the following occurs:

- One row is added to the Survey table.
- One row is added to the Response table for each question in the survey.
- The Question table contains the text of each survey question. The third question in each survey response is an overall satisfaction score. Customers can submit a survey after each purchase.

User Problems

The analytics team has large volumes of data, some of which is semi-structured. The team wants to use Fabric to create a new data store.

Product data is often classified into three pricing groups: high, medium, and low. This logic is implemented in several databases and semantic models, but the logic does NOT always match across implementations.

Requirements

Planned Changes

Litware plans to enable Fabric features in the existing tenant. The analytics team will create a new data store as a proof of concept (PoC). The remaining Litware users will only get access to the Fabric features once the PoC is complete. The PoC will be completed by using a Fabric trial capacity. The following three workspaces will be created:

- AnalyticsPOC: Will contain the data store, semantic models, reports pipelines, dataflow, and notebooks used to populate the data store
- DataEngPOC: Will contain all the pipelines, dataflows, and notebooks used to populate OneLake
- DataSciPOC: Will contain all the notebooks and reports created by the data scientists. The following will be created in the

AnalyticsPOC workspace:

- A data store (type to be decided)
- A custom semantic model
- A default semantic model

Interactive reports

The data engineers will create data pipelines to load data to OneLake either hourly or daily depending on the data source. The analytics engineers will create processes to ingest, transform, and load the data to the data store in the AnalyticsPOC workspace daily. Whenever possible, the data engineers will use low-code tools for data ingestion. The choice of which data cleansing and transformation tools to use will be at the data engineers' discretion.

All the semantic models and reports in the Analytics POC workspace will use the data store as the sole data source.

Technical Requirements

The data store must support the following:

- Read access by using T-SQL or Python
- Semi-structured and unstructured data
- Row-level security (RLS) for users executing T-SQL queries

Files loaded by the data engineers to OneLake will be stored in the Parquet format and will meet Delta Lake specifications.

Data will be loaded without transformation in one area of the AnalyticsPOC data store. The data will then be cleansed, merged, and transformed into a dimensional model. The data load process must ensure that the raw and cleansed data is updated completely before populating the dimensional model. The dimensional model must contain a date dimension. There is no existing data source for the date dimension. The Litware fiscal year matches the calendar year. The date dimension must always contain dates from 2010 through the end of the current year.

The product pricing group logic must be maintained by the analytics engineers in a single location. The pricing group data must be made available in the data store for T-SOL queries and in the default semantic model. The following logic must be used:

- List prices that are less than or equal to 50 are in the low pricing group.
- List prices that are greater than 50 and less than or equal to 1,000 are in the medium pricing group.
- List prices that are greater than 1,000 are in the high pricing group.

Security Requirements

Only Fabric administrators and the analytics team must be able to see the Fabric items created as part of the PoC.

Litware identifies the following security requirements for the Fabric items in the AnalyticsPOC workspace:

- Fabric administrators will be the workspace administrators.
- The data engineers must be able to read from and write to the data store. No access must be granted to datasets or reports.
- The analytics engineers must be able to read from, write to, and create schemas in the data store. They also must be able to create and share semantic models with the data analysts and view and modify all reports in the workspace.
- The data scientists must be able to read from the data store, but not write to it. They will access the data by using a Spark notebook
- The data analysts must have read access to only the dimensional model objects in the data store. They also must have access to create Power BI reports by using the semantic models created by the analytics engineers.
- The date dimension must be available to all users of the data store.
- The principle of least privilege must be followed.

Both the default and custom semantic models must include only tables or views from the dimensional model in the data store. Litware already has the following Microsoft Entra security groups:

FabricAdmins: Fabric administrators

- AnalyticsTeam: All the members of the analytics team
- DataAnalysts: The data analysts on the analytics team
- DataScientists: The data scientists on the analytics team
- DataEngineers: The data engineers on the analytics team
- AnalyticsEngineers: The analytics engineers on the analytics team

Report Requirements

The data analysts must create a customer satisfaction report that meets the following requirements:

- Enables a user to select a product to filter customer survey responses to only those who have purchased that product.
- Displays the average overall satisfaction score of all the surveys submitted during the last 12 months up to a selected date.
- Shows data as soon as the data is updated in the data store.
- Ensures that the report and the semantic model only contain data from the current and previous year.
- Ensures that the report respects any table-level security specified in the source data store.
- Minimizes the execution time of report queries.

You need to ensure the data loading activities in the AnalyticsPOC workspace are executed in the appropriate sequence. The solution must meet the technical requirements.

What should you do?

- A. Create a dataflow that has multiple steps and schedule the dataflow.
- B. Create and schedule a Spark job definition.
- **C. Create a pipeline that has dependencies between activities and schedule the pipeline.**
- D. Create and schedule a Spark notebook.

Answer: C

Explanation:

Pipeline can ensure the activities follow the required sequence.

<https://learn.microsoft.com/en-us/fabric/data-factory/activity-overview#data-transformation-activities>

NEW QUESTION # 68

Hotspot Question

You have a Fabric tenant.

You need to configure OneLake security for users shown in the following table.

The solution must follow the principle of least privilege.

Which permission should you assign to each user? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Answer:

Explanation:

